Comments from the Review committee:

Frank Chlebana, Eric James, and Jonathan Lewis

Summer 2005 should be the target to have all updates complete.

We thought that the DAQ simulation should be updated with the current operating parameters and assumptions. In order to better understand the expected performance of either a new TDC or silicon DAQ and SVT upgrades, we need to use the complete Yale simulation with timing info based on hit distributions in the data. We also need to include into the simulation the effect of the pulsar sending data to 4 CPUs for L2 processing.

At least for internal consumption, it would be interesting to come up with a realistic high luminosity trigger table that includes b-physics.

EVB

This upgrade is well motivated, and we did not have any major concerns about this. The schedule also looks reasonable.

The EVB group should optimize the arrangement of VRBs remaining after COT moves to the new readout scheme in order to evenly spread the data volume load. The optimization ought to include the notion of readout lists. We don't pay a penalty for the XFT and XTRP data very often. We might also consider dropping the DCAS info as well. But mostly, the biggest gains will probably come from dividing up the ISL and L00 more finely.

Does the TDC readout upgrade need to be installed with the upgraded event builder or can this be staged in later?

Pulsar

This upgrade is well motivated and does look like it is progressing well.

Based on the current progress of this group, it seems likely that next summer's shutdown is a realistic target for the L2 upgrade. We should definitely push towards this goal as hard as possible. If we wind up doing numerous major upgrades in 2005, it will be even harder to get back to normal running after that shutdown.

With this in mind, it would be good to try to have the muon trigger part working when we come up from this shutdown.

We did not see any projections of the expected performance. Is the SVT is the limiting factor? Would like to characterize the performance of the L2 upgrade in order to have a system wide simulation.

Assuming that the L2 upgrade happens in summer 2004, we will need to get VRB pulsar readout working first since this comes before the upgraded EVB. Once this happens, will there be any advantage to switching the readout to Gigabit Ethernet?

Need to stage in during a shutdown so coming up with an installation schedule is critical.

Need to test splitting RECES signals. What is minimum number of splitters to test sensibly? We have concerns about the robustness of those fibers.

Silicon

We did not think that this upgrade was well motivated and that there was insufficient justificaion for it. We did not get an impression of what the current L1 rate capability is with the 2 SRC option and what rates one can achieve with the existing system if additional work is done. We need to see ModSim studies that include pulsar and SVT improvements in order to get an idea of the expected performance for both the current system and the upgraded system.

The optimized readout also assumes that noisy channels are kept under strict control. The extent that this can be done and the impact if it is not should be characterized.

Talking with silicon experts afterward it sounded like they would be able to achieve 50KHz at L1 with the existing Silicon with additional work.

It would be good to have Run IIb trigger tables defined that include the anticipated new b triggers.

The option of using the pulsar in place of the SRC should be evaluated. We've not seen any available manpower allocated to bring up a new board and then build its firmware. With the pulsar, they would be starting from a working board and established firmware infrastructure.

Separate clock frequencies for digitalization and readout on the SVX chips could cause significant problems and the reported time gain (some fraction of 2us) seems relatively minor in comparison.

SVT

Road Warrior and AM upgrades seem well-motivated.

Need to understand if there are additional benefits to be gained from a new track fitter board.

What is the available manpower for both the AM and system upgrades?

What is the expected timing performance? We need to know the expected performance of both the 128k-pattern AM upgrade and the final Road Warrior implementation to feed back into a system performance simulation.

Justification for 1M pattern lookup table was lacking. That will need to be motivated better to be considered.

TDC

If the goal is to achive a 1KHz L2 accept rate then we clearly will need to modify the TDC readout and possibly replace the TDCs.

Achieving the 1KHz rate with the existing TDCs would be preferable. In Myron's talk he said that he thinks 1KHz is atainable, but he did not have measurements to back it up.

In light of this we think that the TDC upgrade project should be pursued with continuing R&D on the new TDCs in parallel with the development of the mezzanine card approach. At some point we can then re-evaluate the options and make the selection when both options are more fully developed.

Not having to replace the 9U boards would most likely be preferable. However, we need a demonstration of both concepts to make a choice. In either case, we also need to show that we will be able to tolerate the necessary L1A rate at high luminosity that would lead to a L2A rate as large as 1khz.

It would be nice to be able to combine the readout scheme for the old and the new TDCs, but it looks like the front end part is very different. It is possible that one can try to have a common backend scheme so that one does not have to duplicate efforts, i.e. concentrating the ethernet lines into a switch and have the EVB handle the data fragments in a universal way.

It would be preferable to limit the number of technologies used in the 2b upgrade and have a crate controller with GigE onboard rather than a separate PC with PCI to VME interface. It appears that there are several cards available with this option.

There was also talk about reading the pulsar out in a similar fashion. It is important that there is some general framework/guideline that should be followed so that we can try to make the readout as similar as possible for different systems. The people involved in the EVB upgrade should be fully involved in developing this approach.

Ting mentioned that there were limitations based on the existing Cypress VME interface chip on TDC which slows the VME readout. We need to better understand exactly what the rate limitations are on VME readout over the backplane.

In the meantime, We need to agressively pursue the incremental upgrades to the existing TDCs in order to ensure that the latency remains under control. In particular, we need to continue to pursue the proposed changes to the data format.

XFT

It appears that it will be difficult to meet the schedule. That may be a less serious concern if the new boards are plug compatable with the old, and we don't need to make changes in the hall, i.e. reprogramming XTC cards. However, commissioning with beam would still be unpleasant if the system arrives late.

How is the XFT plan tied to the choice for TDCs and what impact does this have on the design of the hardware? Do we need to make a decision on the TDC option soon or can that be postponed without impacting the XFT upgrade?

With this in mind, it's important to determine as soon as possible whether the exisiting mezzanine cards can produce the required six time bins. The design and fabrication of several hundred new mezzanine cards could be limiting factor on when the upgrade can be completed.

The effect of degraded resolution for the existing system at high luminosity on the SVT performance also needs to be studied.

Is it sufficient just to do the 3D upgrade, i.e. ask for a segment of consistent pt in the window for high-pt tracks? The 3D 2-track trigger gives new functionality.

It's important to think about how this upgrade could be implemented in stages. For example, Richard noted that it would be possible to upgrade the Finders first and the Linkers later. A prioritized list of staged upgrades would enable us to spread out commissioning and help to ensure that a substantial fraction of the new system would be in and operational by the end of the summer 2005 shutdown.

Additional Comments

Jim Patrick

Due to being at a conference out of the country, I was unable to attend the actual review sessions. Below are some comments based on reading the slides and review report. Apologies for any misunderstanding of what was presented.

J. Patrick

Overall I think the report is pretty good and I agree with most of it and don't have a lot to add.

The EVB and L2 upgrades appear to be in good shape. I think they should both be pushed aggressively, a worry is data sizes will go up with luminosity and things may degrade too much with the current systems.

The SVT upgrade looked reasonably well thought out also, I don't have anything to add to that or the XFT parts.

Though potentially interesting, the SVX DAQ upgrade proposal did not seem fully developed. Given past history it would seem optimistic to have a new SRC and FTM boards in only 6 months. There would need to be a firm case based on system modeling that the upgrade would be useful, and also a firm plan and commitments from people to execute it.

A worry that didn't seem to get mentioned is how is the SVX occupancy expected to evolve as the signal/noise deteriorates over time and what effect that has on the overall system performance.

Regarding the proposed TDC readout scheme presented in Ting's slides, qualitatively I do believe in migrating to PC based architectures running Linux. However one should carefully consider if it is worth the effort to develop and support unique technologies in the COT readout. The gain in readout speed compared to the tracer link is modest (~25%) and almost certainly less than alternatives using ethernet. In addition to event readout, there is a lot of other of other software in TDCtest and the diagnostic framework, track test, calibration, and the general front-end framework to interface to the data acquisition control and monitoring programs. I would not recommend removing the current PPC CPUs to run the old TDCs unless the new TDCs are not built, there would be a lot of throwaway software porting involved. Also, it is not obvious why a custom interface to Level 3 would be needed, it would seem gigabit ethernet would be more than adequate. I would suggest evaluating gigabit ethernet throughput on the new 450 MHz 2400 PPC processors. Crate processors could send the data to EVB Scanner CPUs by ethernet, this is equivalent to what D0 does. The rate would likely be fast enough for the needs here. If 2400's aren't fast enough, either an MVME5500 or the Intel EVB board would be, though both would require buying a new version of VxWorks and have support issues. Anyway the point is just to remember that these processors will require

substantially more software and support effort than just a simple event readout routine, and also will require someone to look after them as long as the experiment runs. If a major change in the readout architecture is to be made, it should provide compelling performance and/or maintenance advantages. I don't think there is a performance advantage at all, and given the other 85% of the system will continue to use the existing VxWorks/tracer link system it would seem to increase the maintenance burden.

Regarding the proposal to install Atlas-style ethernet mezzanine cards on the current TDCs: With this sort of architecture with many ethernet sources care must be taken in switch selection, network design, and data transfer protocol to ensure packets are not lost. Otherwise the throughput can be very low. Such a scheme could likely be made to work provided there is sufficient buffering on the TDCs. However it would require some R&D as well as close coordination with the EVB project.

Regarding this quote in the review report:
> > Ting mentioned that there were limitations based on the existing Cypress
> > VME interface chip on TDC which slows the VME readout. We need to better
> > understand exactly what the rate limitations are on VME readout over the
> > backplane.

I believe the Cypress chips are actually fairly fast as they prefetch the data. The rate limit is due to the VME processors, which aren't very fast in 32 bit mode though are >2x as fast in 64 bit mode, which the old TDCs don't support. The problem is this data prefetch causes troubles as the TDC readout is from FIFOs, so the DSP code ensures only one event is in the FIFO at a time. I believe a solution to this ("dual port FIFO mode") is in hand but not yet implemented.